

AMENDMENTS TO THE CLAIMS

This listing of claims replaces all prior versions and listings of claims in the application.

Please cancel claims 5, 9, and 14 without prejudice.

Listing of Claims:

1. (currently amended) A computer-implemented method for identifying optimal allocations of computing resources in a data processing arrangement having a plurality of computing machines that host a plurality of application processes, comprising:

establishing a plurality of server models, each server model including one or more server nodes, wherein each server node represents hardware and has an associated set of capacity attributes;

designating a layered relationship between the server models, wherein for a first server-model layer immediately above a second server-model layer, the second server-model layer includes respective models that represent the nodes in the first server-model layer;

establishing a plurality of service models, each service model including one or more service nodes, wherein each service node represents software and has an associated set of demand attributes and an associated set of capacity attributes;

designating a layered relationship between the service models, wherein for a first service-model layer immediately above a second service-model layer, the second service-model layer includes respective models that represent the nodes in the first server-model layer, wherein a first service node in the first service-model layer is mappable to at least two or more service nodes in the second service-model layer;

normalizing the capacity attributes and the demand attributes; ~~and~~
generating a mapping of service nodes in a first service model in the first service-model layer to service nodes in a second service model in the second service-model layer as a function of the demand attributes of service nodes of the first service model and capacity attributes of service nodes the second service model; and

generating an optimized mapping of service nodes in a user-selected service model to server nodes in a user-selected server model as a function of the associated sets of demand and capacity attributes.

2. (original) The method of claim 1, wherein the capacity attributes include processing capacity and storage capacity and the demand attributes include processing demand and storage demand.

3. (original) The method of claim 2, further comprising:

establishing one or more service-node relationships between selected pairs of the service nodes, wherein each service-node relationship has an associated transport demand attribute specifying a quantity of communication resources required for communication between the associated pair of service nodes;

establishing one or more server-node relationships between selected pairs of the server nodes, wherein each server-node relationship has an associated transport capacity attribute specifying a quantity of communication resources available for communication between the associated pair of server nodes;

normalizing the transport demand attributes; and

generating the optimized mapping as a function of the service-node relationships and server-node relationships.

4. (original) The method of claim 3, further comprising:

selecting one of the server nodes as a base server node;

normalizing the capacity attributes relative to the capacity attribute of the base server node;

selecting one of the service nodes as a base service node; and

normalizing the demand attributes relative to the demand attribute of the base service node.

5. (canceled)

6. (currently amended) The method of claim 4 [[5]], wherein each server node has an associated set of demand attributes and further comprising generating an

optimized mapping of server nodes in a first user-selected server model to server nodes in a second user-selected server model as a function of the demand attributes of the first server model and capacity attributes of the second server model.

7. (original) The method of claim 1, further comprising:

- establishing one or more service-node relationships between selected pairs of the service nodes, wherein each service-node relationship has an associated transport demand attribute specifying a quantity of communication resources required for communication between the associated pair of service nodes;

- establishing one or more server-node relationships between selected pairs of the server nodes, wherein each server-node relationship has an associated transport capacity attribute specifying a quantity of communication resources available for communication between the associated pair of server nodes;

- normalizing the transport demand attributes; and

- generating the optimized mapping as a function of the service-node relationships and server-node relationships.

8. (original) The method of claim 1, further comprising:

- selecting one of the server nodes as a base server node;

- normalizing the capacity attributes relative to corresponding capacity attributes of the base server node;

- selecting one of the service nodes as a base service node; and

- normalizing the demand attributes relative to corresponding demand attributes of the base service node.

9. (canceled)

10. (original) The method of claim 1, wherein each server node has an associated set of demand attributes and further comprising generating an optimized mapping of server nodes in a first user-selected server model to server nodes in a second user-selected server model as a function of the demand attributes of the first server model and capacity attributes of the second server model.

11. (currently amended) An apparatus for identifying optimal allocations of computing resources in a data processing arrangement having a plurality of computing machines that host a plurality of application processes, comprising:

means for establishing a plurality of server models, each server model including one or more server nodes, wherein each server node represents hardware and has an associated set of capacity attributes;

means for designating a layered relationship between the server models, wherein for a first server-model layer immediately above a second server-model layer, the second server-model layer includes respective models that represent the nodes in the first server-model layer;

means for establishing a plurality of service models, each service model including one or more service nodes, wherein each service node represents software and has an associated set of demand attributes and an associated set of capacity attributes;

means for designating a layered relationship between the service models, wherein for a first service-model layer immediately above a second service-model layer, the second service-model layer includes respective models that represent the nodes in the first server-model layer, wherein a first service node in the first service-model layer is mappable to at least two or more service nodes in the second service-model layer;

means for normalizing the capacity attributes and the demand attributes; ~~and~~

means for generating a mapping of service nodes in a first service model in the first service-model layer to service nodes in a second service model in the second service-model layer as a function of the demand attributes of service nodes of the first service model and capacity attributes of service nodes the second service model; and

means for generating an optimized mapping of service nodes in a user-selected service model to server nodes in a user-selected server model as a function of the associated sets of demand and capacity attributes.

12. (currently amended) A system for identifying optimal allocations of computing resources in a data processing arrangement having a plurality of computing machines that host a plurality of application processes, comprising:

a model repository including a plurality of server models and a plurality of service models, each server model including one or more server nodes and each server node representing hardware and having an associated set of normalized capacity attributes, each service model including one or more service nodes and each service node representing software and having an associated set of normalized demand attributes and an associated set of normalized capacity attributes, wherein the server models are defined in a layered relationship and for a first server-model layer immediately above a second server-model layer, the second server-model layer includes respective models that represent the nodes in the first server-model layer, and the service models are defined in a layered relationship and for a first service-model layer immediately above a second service-model layer, the second service-model layer includes respective models that represent the nodes in the first service-model layer, wherein a first service node in the first service-model layer is mappable to at least two or more service nodes in the second service-model layer; and

an optimization engine coupled to the model repository, the optimization engine configured to generate a mapping of service nodes in a first service model in the first service-model layer to service nodes in a second service model in the second service-model layer as a function of the demand attributes of service nodes of the first service model and capacity attributes of service nodes the second service model, and generate an optimized mapping of service nodes in a user-selected service model to server nodes in a user-selected server model as a function of the associated normalized demand and capacity attributes.

13. (original) The system of claim 12, further comprising:

wherein the model repository further includes one or more service-node relationships between selected pairs of the service nodes, each service-node relationship having an associated transport demand attribute that specifies a normalized quantity of communication resources required for communication between the associated pair of service nodes;

wherein the model repository further includes one or more server-node relationships between selected pairs of the server nodes, each server-node relationship having an associated transport capacity attribute that specifies a normalized quantity of communication resources available for communication between the associated pair of server nodes; and

the optimization engine is further configured to generate the optimized mapping as a function of the service-node relationships and server-node relationships.

14. (canceled)

15. (original) The system of claim 12, wherein each server node has an associated set of normalized demand attributes and the optimization engine is further configured to generate an optimized mapping of server nodes in a first user-selected server model to server nodes in a second user-selected server model as a function of the demand attributes of the first server model and capacity attributes of the second server model.

16. (new) The method of claim 1, further comprising:

wherein the server-model layers include a data-center layer, a cluster layer, and a servers layer, the data-center layer having a data-center model representing a plurality of sites, each site having at least one server, the cluster layer having at least one cluster model representing a set of servers, the servers in a set being communicatively coupled, and the servers layer having at least one servers model representing servers in the data processing arrangement;

wherein the service-model layers include a distributed-applications layer, an application-tasks layer, and an application-process layer, the distributed-applications layer having at least one distributed-application model representing a set of application tasks that are mappable to the plurality of sites in the data-center layer, the application-tasks layer having at least one application-tasks model representing tasks that are mappable to a set of servers in the cluster layer, and the application-process layer having at least one application model representing processes that are mappable to servers in the servers layer;

comparing capacity attributes of the at least one data-center model to demand attributes of the at least one distributed-application model; and

mapping application tasks from the at least one distributed-application model to sites in the at least one data-center model based on comparison of the capacity attributes of the at least one data-center model to demand-attributes of the at least one distributed-application model.

17. (new) The method of claim 16, further comprising:

wherein the server-model layers further include a virtual-service-center layer having at least one virtual-service-center model representing a set of sites from the data-center model;

wherein the service-model layers further include a services layer having at least one service model representing a set of application tasks from a distributed-application model;

comparing capacity attributes of the at least one virtual-service-center model to demand attributes of the at least one service model in the services layer; and

mapping application tasks from the at least one virtual-service-center model to the sites in the at least one virtual-service-center model based on comparison of the capacity attributes of the at least one virtual-service-center model to demand attributes of the at least one service model of the services layer.

18. (new) The method of claim 1, further comprising:

wherein the server-model layers include a data-center layer, a cluster layer, and a servers layer, the data-center layer having a data-center model representing a plurality of sites, each site having at least one server, the cluster layer having at least one cluster model representing a set of servers, the servers in a set being communicatively coupled, and the servers layer having at least one servers model representing servers in the data processing arrangement;

wherein the service-model layers include a distributed-applications layer, an application-tasks layer, and an application-process layer, the distributed-applications layer having at least one distributed-application model representing a set of application tasks that are mappable to the plurality of sites in the data-center layer, the application-tasks layer having at least one application-tasks model representing

tasks that are mappable to a set of servers in the cluster layer, and the application-process layer having at least one application model representing processes that are mappable to servers in the servers layer;

comparing capacity attributes of the at least one application model to demand attributes of the at least one application-tasks model in the services layer; and

mapping tasks from the at least one application-tasks model to processes of the at least one application model based on comparison of the capacity attributes of the at least one application model to demand attributes of the at least one application-tasks model.